Attachment A

Application Cover Sheet

Name of Respondent Dell Marketing, L.P.			
Mailing Address One Dell Way	City/Town Round Rock	State Texas	Zip Code 78682
Telephone 617-955-7355	Fax n/a	Web Address www.dell.com	
Primary Contact for Clarificatio Sean Tynan		Primary Contact E-mail Address Sean.Tynan@Dell.com	
Authorized Signatory Laura Valdes		Authorized Signatory E-mail Address Laura.Valdes@Dell.com	
Legal Status/Jurisdiction(e.g., a Massachusetts Corporation, LLC, LLP, etc.)		Respondents UEI No. HJLSY7RZV8X6	

Limited Partnership

Proposal for Massachusetts Technology Collaborative

Request for Proposals for AI Hub Data Commons Collaborative Consulting Services RFP No. 2026-JAII-01

Due: October 7, 2025 by 5:00PM EST

DELL Technologies

One Dell Way
Round Rock, TX 78682
USA
www.dell.com

October 6, 2025

Sabrina Mansur
Procurement Team Leader
Massachusetts Technology Collaborative
75 North Drive
Westborough, MA 01581

Dear Ms. Mansur,

Thank you for this opportunity to submit a proposal for Massachusetts Technology Collaborative's forthcoming technology project. We have studied the information provided to us about your business requirements and carefully analyzed your technology needs. The solution recommended for MassTech has been designed to meet your needs in the most cost-effective way without compromising on quality, service or ongoing support.

Dell is helping our customers to bring down the Total Cost of Ownership by simplifying IT. We are committed to providing solutions that will allow MassTech to reclaim time and cost and increase the productivity of your IT. In addition, we have built environmental consideration into every stage of the Dell product lifecycle including power consumption, helping our customers demonstrate environmentally responsible procurement.

Along with award winning products and services, Dell also offers you a dedicated program account team that is committed to working with you and your procurement needs. This team includes:

- An Account Manager to ensure overall account satisfaction.
- System Consultants to provide a seamless deployment experience.
- Technical Sales Representatives to facilitate order management.
- Customer Service Representatives to provide post-sale support.

Dell looks forward to working with you on this project. Should you have any questions regarding this response, please contact your Account Executive, Sean Tynan at 617-955-7355 or online at Sean.Tynan@Dell.com .

Sincerely,

Laura Valdes
Proposal Manager

Table of Contents

Contacts and Confidentiality Statement 4
Executive Summary 5
Services Required 6
Application 34
Attachment B – Authorized Respondent's Signature and Acceptance Form 36
Attachment C – Budget Template 37
W-9 38
Proposal Legal Notes 39

Contacts and Confidentiality Statement

Contacts

Name	Responsibility	Contact Information
Art Brieva	Services Account Executive	516-526-1254 Art.Brieva@Dell.com
Sean Tynan	Account Executive	617-955-7355 Sean.Tynan@Dell.com
Laura Valdes	Proposal Manager	512-741-6864 Laura.Valdes@Dell.com

Confidentiality Statement:

Dell, and other trademarks are trademarks of Dell Inc. or its subsidiaries.

All other trademarks used herein are the property of their respective owners.

© Copyright 2025 Dell Inc. or its subsidiaries. All Rights Reserved.

This proposal contains confidential material that is proprietary to Dell Technologies Inc. The materials, ideas and concepts contained herein are to be used exclusively to evaluate the capabilities of Dell Technologies Inc. to provide a solution for Massachusetts Technology Collaborative. The information and ideas herein may not be disclosed to anyone outside of Massachusetts Technology Collaborative to be used for purposes other than the evaluation of the capabilities of Dell Technologies Inc. Dell Technologies Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Executive Summary

The Massachusetts AI Hub Data Commons Collaborative (DCC) represents a visionary initiative to create a foundational data ecosystem for AI innovation across the Commonwealth. The Request for Proposal (RFP No. 2026-JAII-01) outlines a complex and multifaceted set of requirements that demand more than a single, monolithic platform; it requires a thoughtfully architected solution that balances enterprise-grade performance, security, and reliability with flexibility, transparency, and community-driven ethos of open-source technologies.

The proposed Dell solution is engineered precisely to meet this dual mandate. At its core are the **Dell Data Lakehouse (DDL)** and **Dell AI Factory (DAF)** platforms. The **Dell Data Lakehouse (DDL)** is a turnkey modern data platform that provides the robust, scalable, and secure foundation necessary for a state-level initiative. The DDL's architecture, built on industry-leading Dell PowerEdge compute and Dell ECS/ObjectScale object storage, and powered by the high-performance Dell Data Analytics Engine (Starburst) and Dell Data Processing Engine (Apache Spark), delivers the raw power to manage, process, and analyze petabyte-scale datasets. The DAF/DDL are sized to customer specifications based on performance metrics that include size of data stored on the platform and AI workloads that would be applied to the platform. None of those parameters were provided in the RFP or the Q&A responses. In good faith, pricing for the hardware will be provided during solution scoping upon down selection.

The **Dell AI Factory (DAF)** is an enterprise-ready solution that helps organizations rapidly adopt and scale artificial intelligence by combining Dell's AI-optimized infrastructure with NVIDIA's leading-edge software and accelerated computing. The DAF integrates compute, storage, networking, and professional services to simplify AI deployment, improve productivity, and enhance security and compliance. This approach enables efficient processing of enterprise data, delivers validated solutions tailored to your unique business needs, and supports a broad range of AI use cases. Customers benefit from faster time-to-value, substantial return on investment, and long-term cost efficiencies compared to public cloud or DIY alternatives. The DAF and DDL platforms are on average 70% more cost effective in year one and two than comparable public cloud solutions and almost all large corporate DIY AI solutions.

This enterprise foundation is strategically augmented with a curated selection of best-of-breed, open-source platforms to deliver the specific functional capabilities required by the DCC. This hybrid approach provides the Commonwealth with the best of both worlds: the long-term stability and support of a Tier 1 technology partner in Dell, and the open, extensible, and community-supported functionality required for a public-facing data commons platform.

Services Required

The selected Respondent will provide full-service project management and technical consulting services for the development and delivery of the DCC. The work includes, but is not limited to ("Services"):

1. Program Management & Coordination
    - Develop and maintain an integrated project plan and timeline
    - Conduct regular meetings and stakeholder updates
    - Ensure milestones defined in the DCC project plan are met
2. Technical Architecture & Platform Development
    - Design and build the data commons website and backend infrastructure
    - Implement metadata management, cataloging, and discovery tools (e.g., CKAN, Dataverse)
    - Develop role-based access and identity management system
    - Develop a central ledger for tracking "Commons Credits"
    - Enable API-based data access and usage analytics
3. Dataset Curation and Integration
    - Assist in curating the initial set of datasets
    - Ensure metadata completeness and FAIR data principles
    - Support ingestion pipelines, version control, and provenance tagging
4. Synthetic Data Capability
    - Integrate synthetic data tools (e.g., SDV, Synthea, CARLA)
    - Support synthetic dataset generation and validation pipelines
    - Ensure compliance with privacy and audit policies
5. AI Fairness and Bias Framework
    - Implement bias detection, auditing, and mitigation tools (e.g., AIF360, Fairlearn)
    - Establish bias impact documentation workflows and dashboards
6. Security, Governance & Compliance
    - Define and implement RBAC, MFA, and dataset tiering
    - Ensure logging, encryption, and audit trails
    - Define terms of use and enforce ethical data handling
7. Testing, Rollout, and Adoption
    - Conduct QA across platform functions

- Train end users and administrators
- Prepare documentation and support materials
- Support public launch

8. Deliverables
    - Detailed Project Plan and Timeline
    - Detailed Platform Architecture Documentation
    - Initial Data Catalog and Metadata System
    - Synthetic Data Tooling and Sample Outputs
    - Fairness and Bias Reporting Dashboard
    - Security and Governance Implementation
    - Training and Launch Materials
    - Sample Commons Credits ledger output
    - Explanation of planned industry engagement
    - Pilot user feedback

Integrated Solution Vision

The success of the DCC hinges not on a collection of disparate tools, but on their seamless integration into a single, cohesive platform. The proposed solution is delivered with this principle at its forefront. **Dell Professional Services** will act as the master systems integrator and single point of accountability for the entire project lifecycle. This services-led approach de-risks the project for the Commonwealth by ensuring that all components—from the foundational hardware to the user-facing data portal and specialized AI fairness toolkits—are architected, deployed, configured, and managed as a unified system.

Our team of experts will manage the end-to-end delivery, from initial strategy and architecture design (ProConsult) through implementation and integration (ProDeploy, Accelerator Services) to ongoing operational support and user adoption (ProSupport, Education Services). This comprehensive engagement model guarantees that the DCC is not merely a technology deployment, but a fully realized, operational, and sustainable capability that achieves its stated goals of centralizing access to data, fostering collaboration, and ensuring ethical AI development.

High-Level Architecture

The proposed architecture is a multi-layered system designed for security, scalability, and usability.

1. **Foundation Layer**: This layer consists of Dell's secure and high-performance infrastructure. Dell PowerEdge R660 servers provide the compute horsepower, while Dell ECS or ObjectScale provides a massively scalable, S3-compliant object storage repository for the data lake itself.

2. **Data Processing & Analytics Layer**: The Dell Data Lakehouse software stack runs on the foundation layer. The Dell Data Analytics Engine, powered by Starburst, enables high-performance, federated SQL queries across all data sources, whether hosted in the Lakehouse or in external systems. The Dell Data Processing Engine, powered by Apache Spark, provides the distributed compute environment for large-scale data transformation, synthetic data generation, and AI/ML model training.
3. **Data Catalog & Portal Layer**: As the primary data catalog user interface for the DCC, this layer will be powered by **CKAN**, the world's leading open-source data portal platform. It will provide all required functionality for data discovery, search, metadata management, and community engagement.
4. **Specialized Services Layer**: This layer consists of integrated, containerized open-source modules to deliver unique DCC requirements. This includes the **AIF360** and **Fairlearn** libraries for the AI Fairness and Bias capability. The **Synthetic Data Vault (SDV)** open-source libraries will be used for the Synthetic Data Generation requirement. All three of these open-source frameworks will be integrated into the DAF running within the DDL's managed Kubernetes environment.
5. **Consumption Layer**: This represents the end-users and systems that will interact with the DCC. This includes researchers and developers using the CKAN web portal, programmatic access via APIs, and direct connections from external analytics tools (e.g., Jupyter notebooks, BI dashboards) to the DDL's query engine.

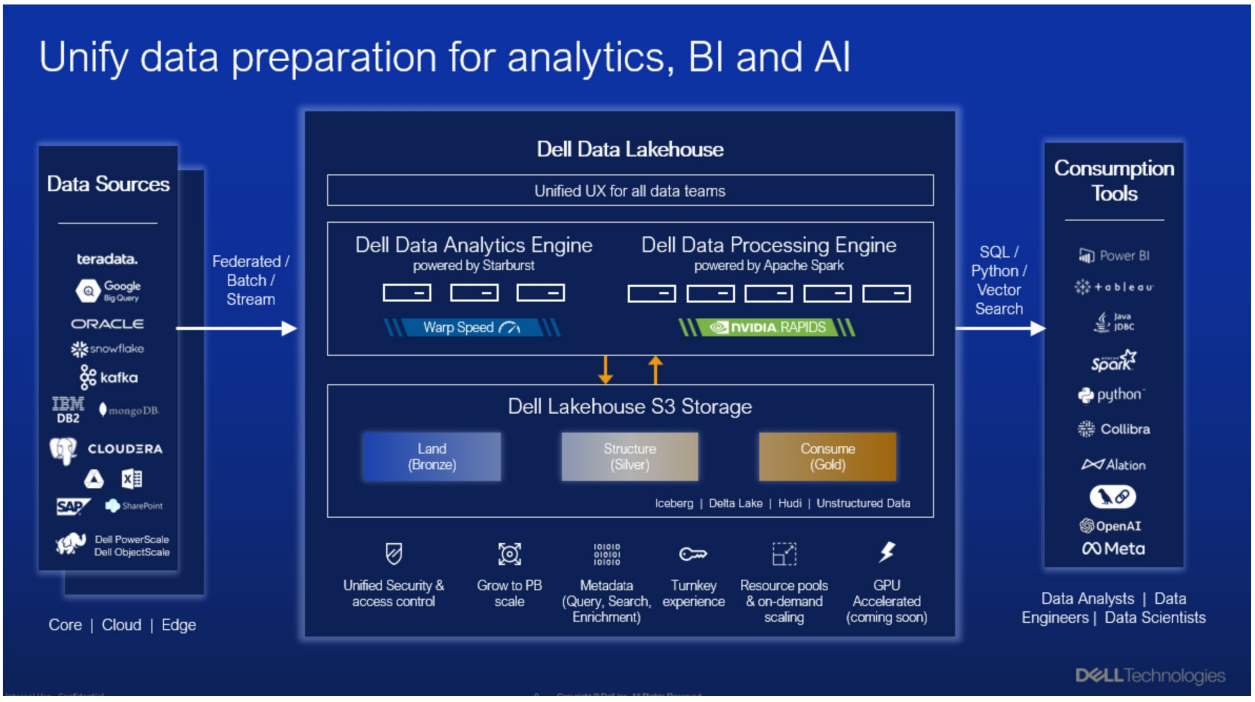

Solution Component Matrix

The following table provides a clear, at-a-glance mapping of the primary RFP requirement areas to the core technical components proposed to address them. This structure demonstrates a purpose-built approach, selecting the best tool for each specific function to create a comprehensive and superior overall solution. This approach also demonstrates the extensibility of the core Dell AI Framework and Dell Data Lake platforms.

RFP Requirement Area	Primary Technology Component	Role in Solution
Website, Dataset Directory, & Data Catalog (Sec 5, 6, 8)	**CKAN** (Open Source)	User-facing portal for data discovery, search, metadata management, and previews.
Hosted Data Platform & Compute (Sec 6, 10, 11)	**Dell Data Lakehouse (DDL)**	Enterprise-grade backend for scalable storage (ECS/ObjectScale), federated query (Starburst), and high-performance AI/ML compute (Spark).
Commons Credits System (Sec 7)	**Custom Development**	A custom development effort to deliver a functional platform that addresses the stated requirements. Solution would combine: • Backend Services: Python (FastAPI) or Go for high-performance, secure APIs. • Database: PostgreSQL. • Asynchronous Tasks: Use a message queue like RabbitMQ for decoupling long-running processes, such as the credit issuance that follows a manual review. • Security: Adhere to OAuth 2.0/OIDC for authentication, implement TLS for all communication, and follow OWASP security best practices.
Synthetic Data Generation (Sec 10)	**Synthetic Data Vault (SDV)**	Integrated Python libraries running on the DDL's Spark engine to generate tabular, time-series, and other synthetic data with Synthetic Data Vault (SDV) framework linked into the DDL.
AI Fairness and Bias Capability (Sec 11, 12)	**AIF360 & Fairlearn** (Open Source)	Integrated Python libraries running on the DDL's Spark engine for dataset auditing, bias detection, and model fairness mitigation.

RFP Requirement Area	Primary Technology Component	Role in Solution
Program Management & Delivery (Sec 2)	**Dell Professional Services**	End-to-end project management, architecture, integration, security implementation, training, and support.

Viability Assessment of Dell Data Lakehouse

The Dell Data Lakehouse (DDL) provides an exceptionally strong and modern foundation for the backend requirements related to hosted data access, management, and interoperability. The platform's architecture directly supports several key tenets of the DCC's mission.

First, the DDL's utilization of open table formats, specifically Apache Iceberg and Delta Lake, natively addresses the critical requirements for robust version control and data lineage tracking. Every change to a dataset is captured as a new version, enabling "time travel" queries that allow researchers to access previous states of the data for reproducibility. This immutable, versioned history is fundamental to the data governance and auditability needs specified in requirements 6.1 and 8.5.

Second, the Dell Data Analytics Engine, powered by Starburst, serves as a powerful, universal query layer. It provides high-performance, standards-based (ANSI SQL) access to all data within the Lakehouse and can federate queries to external data sources without requiring data movement. This directly fulfills the need for programmatic data access via APIs and SDKs, as mandated by requirements 6.1, 8.3, and 8.6. This allows developers and data scientists to connect their preferred tools—from BI dashboards to Python notebooks—directly to the DCC's data assets in a secure, auditable and governed manner.

Finally, the DDL platform is built for scale. Its decoupled compute and storage architecture allows the DCC to independently scale storage capacity (using Dell ECS or ObjectScale) and processing power (by adding PowerEdge compute nodes) as the user base and data volumes grow. This extensibility ensures the platform can support the Commonwealth's ambitions for future expansion without costly re-architecture.

Identified DDL Gap: Data Catalog

While the DDL provides a world-class backend data engine in the form of Starburst, it is not designed as a non-technical user-facing application. It does not include a native, graphical web portal, a browsable dataset directory, or the rich, interactive meta-data management interface required by the entirety of RFP Section 5.0 ("Key Functional Requirements - External Data Sets and Data Capabilities") and significant portions of Sections 6.0 (“Key Functional Requirements - Website Functionality for Hosted Data and User Access”) and 8.0 (“Data Catalog & Metadata Management System”). Fulfilling requirements for a public-facing webpage, dataset detail pages, community submission forms, and an administrative dashboard for content management requires a dedicated front-end portal and data catalog solution.

Proposed Complementary Technology: CKAN

To address the “Identified DDL Gap: Data Catalog”, the proposed solution integrates **CKAN (Comprehensive Knowledge Archive Network)** as the data portal and catalog layer. This choice is directly informed by the RFP, which explicitly suggests CKAN as a potential solution for consideration. Our analysis confirms that CKAN is the de facto global standard for government and research open data portals, powering national sites for the US (data.gov), UK (data.gov.uk), Canada, and Australia, among

others. Its maturity, extensive feature set, and vibrant open-source community make it the lowest-risk and most capable choice for the DCC. Note that the CKAN solution will be wired into the DDL Starburst Galaxy – Gravity component and act as a front end for managing the data catalog. Technical resources will still be able to work directly with Gravity data catalog. Non-technical users will access the CKAN interface.

CKAN natively provides the precise functionality required by the RFP:

- **Dataset Directory and Discovery**: A searchable and filterable directory with customizable "dataset cards," faceted search (by data type, source, domain, license), and full-text search capabilities, directly meeting requirements for the entirety of section 5(Key Functional Requirements – External Data Sets and Data Capabilities) and the entirety of section 8(Data Catalog and Metadata Management System).
- **Rich Metadata Management**: A flexible meta-data schema that supports established standards such as Dublin Core and DCAT, as required by 8.2. It allows for the capture of descriptive, technical, and provenance metadata for each dataset.
- **User-Facing Features**: Functionality for dataset detail pages, community submissions via web forms, and user feedback mechanisms, fulfilling requirements 5.3 and 5.7.
- **Data Previews and Visualization**: An extensible data viewer framework that allows for in-browser previews of common file formats (e.g., CSV, JSON) and can be extended to support more complex visualizations, addressing requirement 8.9.
- **Administration and Governance**: A comprehensive administrative dashboard for managing datasets, users, organizations, and submissions, with role-based access controls, fulfilling requirements 5.8 and 8.4.
- **Extensibility and Interoperability**: A powerful API for programmatic catalog management and a robust plugin architecture. This extensibility is key to our integration strategy, allowing CKAN to connect seamlessly to the DDL's backend storage via existing cloud storage extensions.

Integration Strategy

Integration Strategy: FairLean and AIF360

Dell Professional Services will be responsible for the seamless integration of the Fairlean and the AIF360 bias detection frameworks into the DAF backend. The architecture will operate as follows: Farilean and AIF360 will be deployed within the DDL's managed Kubernetes environment for high availability. Steps will be integrated into the learning model build steps. The Fairlean and AIF360 interfaces will be cataloged and access post-build runs.

Integration Strategy: CKAN Data Catalog

Dell Professional Services will be responsible for the seamless integration of the CKAN front-end with the DDL backend. The architecture will operate as follows: CKAN will be deployed within the DDL's managed Kubernetes environment for high availability. It will serve as the system of record for all metadata, user interactions, and the overall discovery experience.

The data itself will reside on the highly scalable and secure Dell ECS object storage layer of the DDL. Using CKAN's file store and cloud storage extensions, a direct link will be established between the

metadata record in CKAN and the actual data object(s) in ECS. When a user clicks to download a dataset or requests a data preview through the CKAN interface, CKAN will securely broker that request, leveraging the DDL's infrastructure to serve the data. For programmatic access, technical users will discover datasets via the Gravity framework and receive secure, pre-signed URLs to access the data directly from the DDL's high-performance storage or connect to the Starburst API endpoint for direct querying. Both technical and non-technical users will be able view data catalog information based on RBAC configured in the CKAN user interface. The Gravity data catalog component will remain intact and operational to ensure that regular DDL product upgrades and patch releases continue to work Business As Usual under the Dell Managed Service contract and warranty for the DDL platform

This integrated architecture provides a powerful synergy that directly supports the FAIR (Findable, Accessible, Interoperable, Reusable) data principles, which are a cornerstone of the DCC's mission. CKAN excels at making data **Findable** through its rich search and metadata capabilities and **Accessible** through its user-friendly web portal. The DDL ensures data is **Interoperable** by storing it in open formats (e.g., Parquet) accessible via standard APIs (SQL) and **Reusable** through its robust versioning, lineage, and scalability, which are essential for large-scale AI research. This combination creates a platform where each component is purpose-built to fulfill specific pillars of the FAIR framework, resulting in a solution that is far greater than the sum of its parts.

Response to Section 5.0: Key Functional Requirements - External Data Sets and Data Capabilities

Our proposed solution provides a comprehensive, user-friendly, and feature-rich web portal powered by CKAN, the world's leading open-source data portal platform, backed by the scalable and secure Dell Data Lakehouse.

- **Requirement 5.1 Webpage**: The CKAN-based portal will provide a professional and engaging landing page for the Data Commons Collaborative. It will feature an introduction to the DCC's mission, curated sections for highlighted dataset collections (Public, Massachusetts, Synthetic), a prominent quick search bar, and a dynamic section for recent updates or featured datasets, meeting all specified requirements.
- **Requirement 5.2 Dataset Directory**: The solution will deliver a powerful and intuitive dataset directory. Users will be able to perform full-text keyword searches and apply multiple filters simultaneously, including Data Type, Source, Domain, File Format, and License. Each dataset will be represented by a "Dataset Card" containing its name, a brief description, relevant tags, its source/platform, and a direct link to view or download the data, fully satisfying all sub-requirements.
- **Requirement 5.3 Dataset Detail Pages**: Each dataset will have a dedicated detail page that provides comprehensive metadata, including its size, source, full description, license terms, and categories. These pages will also feature links to related datasets and any external documentation, providing researchers with the full context needed to evaluate and use the data.
- **Requirement 5.4 Public Data Sets**: The portal will include a curated section dedicated to high-value public datasets from global repositories like CARLA, CERN, World Bank Open Data, and others listed in Appendix A. Each entry will include a description, license details, and direct download/access links.
- **Requirement 5.5 Synthetic Data Tools Section**: A dedicated section of the portal will serve as a directory for synthetic data tools and simulators like SDV and NVIDIA Omniverse. Each tool will

have a descriptive entry with links to its documentation, source code (e.g., GitHub), and relevant use case examples, as specified.

- **Requirement 5.6 Massachusetts Public Datasets**: The portal will feature a prominent section for datasets from Massachusetts public sources, such as MassGIS and MassDOT. Each dataset will be cataloged with download links, license information, and integration notes to encourage use by civic tech groups and local researchers. The portal section that shows these view only facts will be making API calls to the native Gravity data catalog that is front ended by the CKAN data catalog.
- **Requirement 5.7 Community Dataset Submissions**: The CKAN platform includes built-in functionality for community engagement. We will configure a user-friendly web form that allows users to suggest new datasets or provide feedback on existing ones. This submission and feedback mechanism will be routed to DCC administrators for review, meeting the requirement for community-driven curation.
- **Requirement 5.8 Admin Dashboard**: The solution provides a comprehensive, web-based administrative dashboard within CKAN. This dashboard will allow authorized DCC staff to add, edit, or remove datasets and links; manage and moderate community suggestions; and oversee the health of the catalog. The system supports fine-grained, role-based access controls to ensure proper governance, fulfilling all specified requirements.
- **Requirement 5.9 Additional Requirements**: The proposed solution will meet all additional non-functional requirements.
  - **Accessibility**: The CKAN front-end will be themed and configured to meet the accessibility standards set by MassTech.
  - **Performance**: The combination of the highly optimized CKAN portal and the high-performance DDL backend will ensure page loads and searches are completed under 2 seconds, even with a catalog of over 1,000 datasets.
  - **Security**: The entire platform will be secured with HTTPS to TLS 1.3 standard, and all user inputs will be sanitized. The DDL backend provides robust security, including rate limiting for API access. Dell Managed Services will manage all system patches and currency upgrades to ensure data security and protection is always up to date.
  - **Responsiveness**: The CKAN portal theme will be fully responsive, providing an optimal user experience on mobile, tablet, and desktop devices.
  - **Extensibility**: The solution's architecture is inherently modular, based on CKAN's plugin system and the DDL's API-first design, allowing for the easy addition of future data sources or capabilities.

Response to Section 6.0: Key Functional Requirements - Website Functionality for Hosted Data and Access

Our solution architecture, combining the CKAN portal with the Dell Data Lakehouse backend, provides a robust and scalable platform for hosting and accessing curated datasets.

- **Requirement 6.1 Data Hosting and Access**: The DDL platform provides an enterprise-grade data hosting solution.

- **Metadata Standards**: Metadata will be managed in CKAN, supporting standards like DCAT for discovery and context.
- **Searchable & Downloadable Formats**: Datasets will be searchable and browsable via the CKAN portal and downloadable in machine-readable formats like CSV, JSON, and Parquet, which are natively supported by the DDL's storage layer.
- **APIs and Export Tools**: The solution provides multiple APIs for programmatic access. The CKAN and Starburst Gravity API's allow for programmatic search and discovery of metadata, while the DDL's Starburst query engine provides a powerful SQL API for direct data access and analysis.
- **Version Control and Lineage**: The DDL's use of the Apache Iceberg open table format provides native, transaction-level version control and data lineage for all hosted datasets, ensuring full auditability and reproducibility.

- **Requirement 6.2 User Experience Enhancements**: The CKAN front-end delivers a rich user experience for exploring hosted data.
- **Dataset Explorer**: The portal provides a sophisticated dataset explorer with multi-select filters for category, publisher, update frequency, and other metadata fields.
- **Preview Functionality**: CKAN's data viewer plugins will be configured to provide in-browser previews of hosted data, including sample rows for tabular data and schema views, without requiring a full download.
- **Access Requests**: The platform will include a workflow for requesting access to restricted datasets. Users can submit a request via the portal, which will be routed to administrators for review and approval, with access being granted via the DDL's underlying role-based access control system.

Response to Section 8.0: Data Catalog and Metadata Management System

The proposed solution delivers a world-class data catalog and metadata management system by integrating CKAN as the front-end with the DDL as the backend data platform, fully aligning with the FAIR data principles central to the DCC's mission.

- **Requirement 8.1 Capability Definition and Benefits**: Our solution provides a comprehensive system to organize, describe, and make datasets discoverable and reusable. It fully supports the FAIR principles (Findable, Accessible, Interoperable, Reusable) by combining CKAN's discovery and access features with the DDL's open, interoperable data formats and scalable architecture.
- **Requirement 8.2 Metadata Management**: CKAN provides a rich and flexible metadata management system. It supports descriptive metadata (title, author, keywords, license), technical metadata (format, schema, size), and provenance information. The system is built to support established metadata standards like Dublin Core, Schema.org, and DCAT, ensuring compliance and interoperability.
- **Requirement 8.3 Dataset Discovery and Search**: The solution offers powerful discovery capabilities. CKAN provides faceted search, allowing users to filter by tags, creator, date, and format. Its full-text search engine, powered by Solr, indexes all metadata for fast and relevant

"Google-style" searching. A comprehensive RESTful API enables programmatic search and retrieval of catalog information.

- **Requirement 8.4 Access Control and Governance**: A unified governance model will be implemented across both CKAN and the DDL. This includes role-based permissions (defining who can view, edit, or publish data) and tiered access levels (public, registered, restricted). The DDL provides comprehensive audit trails, logging all data access and modifications for full accountability.
- **Requirement 8.5 Dataset Publishing and Versioning**: The solution provides a streamlined workflow for publishing data. Administrators can use the CKAN interface to easily upload or ingest datasets, which are then stored on the DDL backend. A review and approval workflow can be configured for quality control. Critically, the DDL's use of Apache Iceberg provides immutable version control for all datasets, tracking every change over time.
- **Requirement 8.6 Interoperability and Integration**: The entire platform is built on APIs and open standards. CKAN provides RESTful APIs and SDKs to integrate with external systems. The DDL provides a SQL API (via Starburst) and Spark integration, allowing seamless connection from analytics tools and Jupyter notebooks. The use of open standards ensures FAIR compliance. The system can be extended to support persistent identifiers like DOIs through available CKAN plugins.
- **Requirement 8.7 Data Usage and Metrics**: The platform will track and report on data usage. The DDL's audit logs will capture all download and access events, providing metrics on which datasets are being used and how often. CKAN can be configured to track citations, linking datasets to publications that use them.
- **Requirement 8.8 Collaboration and Community Features**: CKAN is designed to foster community and collaboration. It includes features for users to leave comments and annotations on datasets. It also supports team-based curation, allowing different groups or organizations to manage their own collections of datasets within the broader commons.
- **Requirement 8.9 Data Previews and Exploration**: The CKAN portal will provide in-browser previews for common data types (CSV, JSON, images), allowing users to explore data without downloading it. The DDL's processing engine can be leveraged to generate data profiles and summary statistics, which can be displayed on the dataset detail pages to aid in data exploration.

Response to Section 7.0: Commons Credits System

Viability Assessment of Dell Data Lakehouse

The Dell Data Lakehouse is a data storage and analytics platform. Its native feature set is focused on data processing, querying, governance, and management at scale. The platform does not include any built-in functionality for a credit-based ledger, rewards system, or transactional economic model as described in the RFP. This requirement represents a complete functional gap in the core DDL product. A successful proposal requires the integration of a custom web interface built by Dell Professional Services to deliver this capability.

Identified Gap & Requirement Analysis

The "Commons Credits System" outlined in RFP Section 7.0 is a highly specific and unique requirement for a data platform. It describes a micro-economy designed to incentivize contributions and manage resource consumption within the DCC ecosystem. The core technical requirements are for a **secure, auditable, centralized ledger** to record all credit transactions (issuance, transfer, redemption). This system must be:

- **Interoperable**: It must integrate with the platform's identity access management (IAM) system (e.g., SSO).
- **Rule-Based**: It must support the programmatic implementation of rules for earning credits (e.g., for data contributions) and spending credits (e.g., for accessing premium data or compute resources).
- **API-Driven**: All components must be modular and accessible via APIs to allow integration with the main data portal and other currently DAF embedded platform services like Run.AI.
- **User-Facing**: It must provide a way for users to view their credit balance and transaction history. This will be addressed in the bespoke Common Credits section of the custom DCC web portal.
- **Extensible**: The RFP explicitly notes that a preference for open-source technologies and future support for blockchain integration is a plus.

Building such a system from scratch may be a complex, time-consuming, and high-risk software development project. The worldclass core extensible DAF and DLH platforms were purpose built to be modified for unique customer use cases such as the ones defined by the DCC requirements document.

Common Credits & Centralized Ledger

1.1 Centralized Ledger and Tracking Infrastructure

This is the foundational component of the system, the source of truth for all credit transactions.

- **Solution**: Use a robust, open-source relational database like **PostgreSQL** to serve as the central ledger. PostgreSQL offers transactional integrity (ACID compliance), strong security features, and excellent auditability capabilities.
- **Approach**:
    - **Database Schema**: Design a simple and effective schema. The core would be an immutable transactions table where records are only ever added, never modified or deleted. Corrections should be handled by issuing new, reversing transactions.
        - accounts table: Stores user/organization ID and current balance.
        - transactions table: Contains transaction_id, from_account_id, to_account_id, amount, timestamp, purpose_tag (e.g., model_training_fee), and metadata (a JSONB field for extra details).
    - **Identity Integration**: Do not store user credentials. Link the accounts table to your existing identity management system using a unique user identifier (UUID) provided by your SSO via

protocols like **OAuth 2.0** or **OIDC**. API interactions would be authenticated using JWT Bearer Tokens.

- **API Layer**: Wrap the database with a dedicated **Ledger Service** (a microservice). All other services must interact with the ledger through this secure API, never directly with the database. This enforces business rules and maintains auditability.

1.2 Earning Commons Credits

This involves the "minting" of new credits into the ecosystem.

- **Solution**: Develop an **Issuance Service** that handles all logic for granting credits. This service would be the single-entry point for creating credits, ensuring consistency.
- **Approach**:
  - **Contributions**: Implement a workflow where contributions (e.g., uploading a dataset) enter a "pending review" state. Once an admin or an automated quality check approves the contribution, it triggers an API call to the Issuance Service to credit the contributor's account.
  - **Grants**: Create a simple administrative UI where authorized grant managers can use a form to issue credits to specific users or organizations. This form would make a secure API call to the Issuance Service.
  - **Purchase**: Integrate with a third-party payment processor like **Stripe**. When a user completes a purchase, Stripe sends a webhook notification to an endpoint in your system, which then securely triggers the Issuance Service to grant the purchased credits.

1.3 Spending Commons Credits

This involves the "burning" or redemption of credits for platform resources.

- **Solution**: Implement a **Billing Middleware** that integrates with an **API Gateway** (like Kong, Tyk, or a cloud-native solution). The gateway sits in front of your premium services (e.g., model APIs, data download endpoints) and intercepts requests.
- **Approach**:
  - **Request Interception**: Before forwarding a user's request to the target service, the API Gateway calls the Billing Middleware.
  - **Transaction Logic**: The middleware performs the following steps:
    1. Identifies the user from their API token.
    2. Determines the cost of the requested action (e.g., "1 GPU hour = 10 credits").
    3. Queries the Ledger Service for the user's current balance.

4. If the balance is sufficient, it instructs the Ledger Service to debit the account and allows the request to proceed. If not, it denies the request with a 402 Payment Required error.

- **Enforcement**: Use a scheduled job (e.g., a cron job) to run daily or weekly to enforce rules like credit expiration by issuing debit transactions for expired credits. Throttling and tiered access can be configured directly in the API Gateway based on user roles or metadata.

1.4 Incentivizing Efficient Usage

This requires a dynamic pricing model rather than fixed costs for all actions.

- **Solution**: Create a **Pricing Engine Service**. Instead of hard-coding costs, the Billing Middleware will query this service to get the price for a specific action at a specific time.
- **Approach**: The Pricing Engine calculates costs based on various factors:
  - **Energy Efficiency**: It can query infrastructure metadata to offer a "discount" for training on more energy-efficient GPUs.
  - **Off-Peak Usage**: Charge fewer credits for jobs run during off-peak hours (e.g., overnight) to balance infrastructure load.
  - **Asset Reuse**: Provide a "rebate" or lower access cost for using popular, well-documented models or datasets, incentivizing the use of shared assets over creating redundant ones.

1.5 Governance and Transparency

This is about providing visibility and control to platform administrators.

- **Solution**: Build a dedicated **Admin Dashboard** web application.
- **Approach**:
  - This dashboard will be a secure, role-based front-end that communicates with the credit system's APIs.
  - **Features**:
    - **Policy Management**: A UI for setting and adjusting the rules for earning credits (e.g., "High-quality documentation upload = 50 credits").
    - **Reporting**: Visualize transaction data with charts and graphs showing credit issuance vs. spending, popular services, and top contributors. Tools like **Metabase** or **Apache Superset** can be embedded for powerful, open-source BI capabilities.
    - **User Management**: An interface for viewing user balances, transaction histories, and setting spending quotas for specific users or organizations.

1.6 User Experience & Access Control

This is the user-facing component where individuals manage their credits.

- **Solution**: Integrate a **"Credit Wallet"** section directly into the main Data Commons user interface.
- **Approach**:
    - **UI Components**: Develop a set of components for the user dashboard:
        - A clear display of the current credit balance.
        - A paginated, searchable transaction history.
        - A "Marketplace" or "Services" page showing all the ways users can spend their credits (e.g., available datasets, models, compute options) with clear pricing.
        - A guide on "How to Earn Credits" that links to the contribution workflows.
    - **Access Control**: Implement Role-Based Access Control (RBAC) at the API level. A standard user can only view their own balance and history, while an organizational admin might be able to view the balances of all members in their group.

1.7 Technical Expectations & Proposed Architecture

This architecture meets the requirements for modularity, security, and extensibility.

- **Architecture**: A **Microservices Architecture** where each component is an independent, containerized application deployed on a platform like Kubernetes.

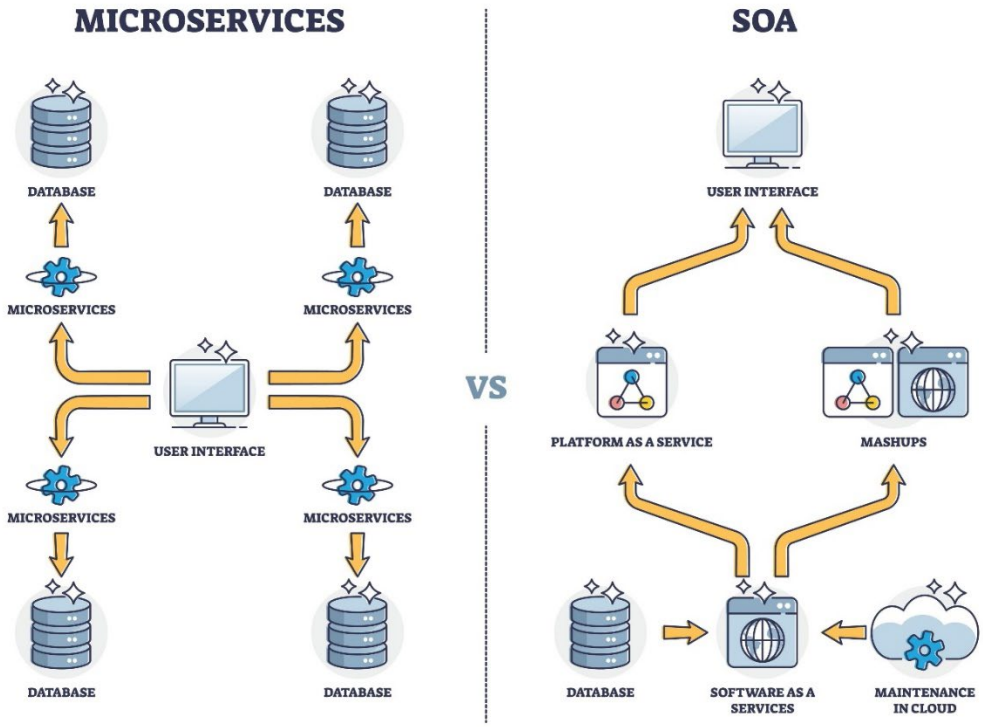

- **Technology Stack**:
  - **Backend Services**: **Python (FastAPI)** or **Go** for high-performance, secure APIs.
  - **Database**: **PostgreSQL**.
  - **Asynchronous Tasks**: Use a message queue like **RabbitMQ** for decoupling long-running processes, such as the credit issuance that follows a manual review.
  - **Security**: Adhere to **OAuth 2.0/OIDC** for authentication, implement TLS for all communication, and follow OWASP security best practices.
- **Future Blockchain Support**: This design is "blockchain-ready." The immutable transaction log in PostgreSQL can be periodically hashed and anchored to a public or private blockchain to create a tamper-evident external audit trail without the complexity and overhead of running a full blockchain from day one.

Building a solution of this scope is a significant undertaking. Below is a detailed breakdown of the estimated resources and a phased timeline.

This estimate assumes a dedicated, experienced team and that certain foundational elements (like a primary cloud provider, SSO, and basic CI/CD pipelines) are already in place.

Key Assumptions

- **Existing Infrastructure**: You have an established cloud environment (AWS, GCP, Azure), a functioning Kubernetes or container orchestration platform, and a mature CI/CD process.
- **Existing Identity Management**: A Single Sign-On (SSO) solution using standards like OIDC/OAuth2 is already in use for the Data Commons platform.
- **Clear Product Ownership**: There is a dedicated product owner who can make decisions quickly and provide clear requirements.
- **Team Availability**: The resource estimates are for a dedicated team focused solely on this project.

---

1. Required Resources (The Project Team)

This is the ideal team composition for building this system efficiently and securely. Some roles could be combined in a smaller team, but this would extend the timeline.

- **Product Manager (1 FTE)**: Owns the vision, prioritizes features, and acts as the liaison between stakeholders and the development team.
- **Tech Lead / Solution Architect (1 FTE)**: Designs the overall system architecture, makes key technology choices, and guides the engineering team.
- **Backend Engineers (2-3 FTEs)**: The core developers responsible for building the microservices (Ledger, Issuance, Billing, etc.), database schema, and APIs. Python/Go experience is ideal.
- **Frontend Engineer (1 FTE)**: Builds the user-facing "Credit Wallet" and the administrative governance dashboard. Experience with a modern framework like React or Vue.js is essential.
- **DevOps / Platform Engineer (1 FTE)**: Manages the infrastructure-as-code, sets up CI/CD pipelines for the new services, and handles deployment, monitoring, and logging.
- **QA Engineer / SDET (1 FTE)**: Develops the testing strategy, creates automated test suites (unit, integration, end-to-end), and ensures the system is reliable and bug-free, especially concerning financial transactions.
- **UI/UX Designer (Part-Time / Contract)**: Needed primarily during the initial design phases for the user wallet and admin dashboard to ensure a high-quality user experience.

Total Team Size: Approximately 8-12 people.

---

2. Phased Development Timeline & Deliverables

We recommend a phased approach to deliver value incrementally and adapt to feedback. The total projected timeline for a full-featured system is **8-10 months**.

Phase 1: Foundation & Core Ledger (Months 1-2)

The goal here is to build the secure, non-negotiable core of the system. This phase will have no user-facing UI.

- **Activities**:
    - Finalize detailed system architecture and technology choices.
    - Design and implement the database schema in PostgreSQL.
    - Build the core **Ledger Service** API (e.g., create accounts, get balance, record transactions).
    - Integrate the service with the existing SSO for authentication.
    - Establish CI/CD pipelines and logging/monitoring for this core service.
- **Key Deliverable**: A functional, secure, and auditable (but "headless") ledger system that can be operated via its API.

Phase 2: Minimum Viable Product (MVP) (Months 3-5)

The goal is to launch the simplest possible end-to-end flow: an admin can grant credits, and a user can spend them on *one* specific service.

- **Activities**:
    - Build the initial **Issuance Service** with a basic function for manual admin grants.
    - Develop the **Billing Middleware** and integrate it with the API Gateway for a single, high-value service (e.g., a specific model inference API).
    - Build a simple **Admin Dashboard** allowing authorized staff to manually issue credits to users.
    - Create the basic **User Wallet UI** where users can view their current balance and a simple transaction history.
- **Key Deliverable**: A live, functioning credit system for a limited set of users and services. This allows for early feedback on the core economic loop.

Phase 3: Feature Enhancement & Automation (Months 6-8)

The goal is to build out the full feature set, focusing on automation and incentives.

- **Activities**:
    - Expand the **Issuance Service** to support automated earning (e.g., API hooks for dataset contributions).
    - Integrate a payment gateway (e.g., Stripe) to allow users to purchase credits.
    - Build the **Pricing Engine Service** to enable dynamic costs for services (e.g., off-peak discounts).
    - Enhance the Admin Dashboard with reporting, analytics, and policy management tools.
    - Expand the Billing Middleware integration to cover all premium platform services.
- **Key Deliverable**: A largely automated and feature-rich credit economy that can scale with the platform.

Phase 4: Governance & Optimization (Months 9-10+)

The goal is to refine the system based on usage data and add advanced governance features.

- **Activities**:
  - Implement credit expiration and throttling rules.
  - Build advanced reports and dashboards for stakeholders on platform sustainability and usage patterns.
  - Fine-tune the incentive mechanisms based on observed user behavior.
  - Perform security audits and performance optimization.
- **Key Deliverable**: A mature, sustainable, and well-governed platform economy.

---

3. Infrastructure & Tooling Resources

These are the non-personnel costs associated with the project.

- **Cloud Hosting**: Costs for databases (e.g., AWS RDS for PostgreSQL), compute instances/containers (Kubernetes cluster), API Gateway usage, logging/monitoring services, and data storage.
- **Software Licensing**: Potential costs for a managed API Gateway, observability platforms (e.g., Datadog, New Relic), and security scanning tools.
- **Payment Gateway Fees**: Transactional fees from the payment provider (e.g., Stripe typically charges ~2.9% + $0.30 per transaction).

Summary Table

Phase	Timeline	Key Deliverables	Resource Focus
1. Foundation	**Months 1-2**	Secure Ledger Service & Database	Backend, DevOps, Tech Lead
2. MVP Launch	**Months 3-5**	Simple end-to-end flow (manual grants, spending on 1 service), basic User/Admin UI	Full Team (Frontend begins)
3. Enhancements	**Months 6-8**	Automated earning, credit purchase, dynamic pricing, full reporting	Full Team
4. Governance	**Months 9-10+**	Advanced rules (expiration), optimization, security audits	Backend, DevOps, Product
4. Warranty Support	**90 Days UAT**	Production support for functional defects discovered after go-live	DUX, Services Developer, DevOps, Database admin

Response to Section 10.0,11.0: Advanced AI and Data Science Capabilities

Viability Assessment of Dell Data Lakehouse

The Dell Data Lakehouse (DDL) platform is an ideal foundation for the advanced AI and data science capabilities required by the DCC. Its architecture is purpose-built to handle the large-scale, computationally intensive workloads associated with modern AI, including synthetic data generation and fairness analysis.

The core of this capability is the **Dell Data Processing Engine**, which is powered by Apache Spark. Apache Spark is the undisputed industry standard for distributed, large-scale data processing and machine learning. By providing a managed, enterprise-grade Spark environment, the DDL enables data scientists and researchers to execute complex tasks on datasets that are far too large to be processed on a single machine. The DCC plans to host multi-terabyte and even petabyte-scale datasets like Common Crawl and MIMIC-IV; applying fairness audits or training generative models on data of this magnitude is only feasible with a powerful distributed engine like Spark.

Furthermore, the DDL platform is designed to provide the necessary infrastructure, including scalable GPU and CPU environments, which are critical for training the generative AI models (e.g., GANs, VAEs) used in synthetic data generation. The platform's containerized execution environments also ensure that data science workflows are reproducible and that the required software libraries and dependencies can be managed effectively.

Identified Gap

While the DDL provides the essential *compute environment* to run these advanced workloads, it does not come pre-packaged with the specific, high-level Python libraries for synthetic data generation or AI fairness that are mentioned in the RFP. The DDL provides the engine, but the "applications" (the specialized toolkits) must be integrated on top of it. This is not a weakness but a reflection of the platform's flexible, open design, which allows for the integration of the best and most current open-source tools.

Proposed Complementary Technologies: SDV, AIF360, and FairLearn

The proposed solution will fully integrate the leading open-source libraries explicitly suggested in the RFP, making them available as first-class tools within the Dell AI Factory and Dell Data Lakehouse platforms.

- **Synthetic Data Generation Capability**: We will integrate a suite of synthetic data generation tools, with a primary focus on the **Synthetic Data Vault (SDV)** library. SDV, developed at MIT, is a powerful open-source library that uses deep learning and probabilistic models (like GANs and VAEs) to generate high-fidelity synthetic tabular and time-series data
- **AI Fairness and Bias Capability**: We will integrate both **AI Fairness 360 (AIF360)** and **Fairlearn**.
    - **AIF360**, an open-source library originally from IBM, provides a comprehensive toolkit with over 70 fairness metrics and a wide array of bias mitigation algorithms. These algorithms cover the entire machine learning lifecycle, including pre-processing (e.g., reweighting data), in-processing (e.g., fairness-aware algorithms), and post-processing (e.g., adjusting model outcomes) techniques. This directly addresses the requirements in Section 12.1 for dataset auditing, model evaluation, and bias mitigation tools.
    - **Fairlearn**, an open-source project with origins at Microsoft, provides a powerful API and interactive dashboard for assessing and comparing model performance across different

subgroups. Its focus on visualization and ease of use makes it an excellent complement to AIF360's algorithmic depth, fulfilling the need for user-friendly tools and reporting dashboards.

Integration Strategy

The key to our strategy is not just to provide these libraries, but to operationalize them at scale. Dell Professional Services will pre-configure the DDL's containerized Spark environments with AIF360, Fairlearn, SDV, and their dependencies. This eliminates the setup and configuration burden for end-users.

Crucially, we will provide a set of template **Jupyter Notebooks** that demonstrate canonical workflows. For instance:

- A notebook will show a researcher how to load a multi-terabyte Parquet dataset from the DDL's ECS storage into a Spark DataFrame, run a comprehensive fairness audit using AIF360's metrics across protected attributes, and visualize the results.
- Another notebook will demonstrate how to use SDV to train a generative model on a sensitive dataset, leveraging the DDL's distributed Spark cluster to parallelize the training process, and then generate a privacy-preserving synthetic version of that dataset.

This approach provides a powerful value proposition. Any data scientist can run these tools on a small CSV file on their laptop. However, our solution enables the DCC's users to apply these critical responsible AI techniques to the massive, real-world datasets that the platform is designed to host. We are delivering not just the tools, but the enterprise-grade, scalable platform required to make them practical and effective.

Response to Section 10.0: Synthetic Data Generation Capability

Our solution provides a complete, end-to-end Synthetic Data Generation Capability, leveraging the powerful Dell Data Processing Engine (powered by Apache Spark) and integrated open-source tools like the Synthetic Data Vault (SDV).

- **Requirement 10.1 Capability Definition and Benefits**: The proposed capability enables the generation of artificial data that mimics the statistical properties of real-world data, enhancing privacy, improving data access, and allowing for the simulation of rare scenarios. This is all accomplished within the secure and scalable DDL environment.
- **Requirement 10.2 Synthetic Data Generation Requirements**:
    - **Data Ingestion & Preprocessing**: The DDL's Spark engine can ingest diverse data types from the data lake storage. We will provide tools and notebooks for automated data profiling and de-identification prior to synthetic processing.
    - **Generation Capabilities**: By integrating SDV and other generative AI libraries, the platform will provide state-of-the-art tools for generating synthetic data. Users can configure synthesis parameters and full traceability and metadata linkage between original and synthetic datasets will be maintained.
    - **Data Storage and Tagging**: Synthetic datasets will be stored in the DDL's object storage layer with clear metadata tags in the CKAN catalog to distinguish them from real data and record their provenance.

- **Compute Infrastructure**: The DDL provides scalable GPU/CPU environments optimized for training generative models and supports containerized execution for reproducible pipelines.
- **Privacy and Compliance**: The platform will enforce all relevant data privacy policies. Integrated tools can be used to validate the privacy guarantees (e.g., differential privacy) of the generated data.
- **Access Control and Governance**: The platform's unified RBAC system will define specific roles and permissions for synthetic data creators and consumers, ensuring proper governance.

Response to Section 11.0: AI Fairness and Bias Capability

The proposed solution establishes a comprehensive AI Fairness and Bias Capability designed to ensure that data and AI systems developed within the DCC are equitable, transparent, and aligned with the highest ethical standards. This capability is powered by the integration of leading open-source toolkits, AIF360 and Fairlearn, running on the scalable DDL platform.

- **Requirement 11.1 Purpose**: Our solution is fundamentally aligned with the guiding principles of being Human-Centric, Trustworthy, Inclusive, Innovative, and Sustainable. By providing practical, scalable tools for identifying and mitigating bias, the platform empowers researchers to build AI systems that serve the public good, foster trust, and promote equitable outcomes for all citizens of the Commonwealth. The integration of AIF360 and Fairlearn directly operationalizes these principles.

Response to Section 12.0,13.0,14.0,2.0: Security, Governance, and Compliance Framework

Viability Assessment of Dell Data Lakehouse

The Dell Data Lakehouse (DDL) platform provides a robust, multi-layered foundation for the comprehensive security, governance, and compliance framework required by the DCC. Security is not an add-on but is engineered into every level of the stack, providing an enterprise-grade posture suitable for managing sensitive public and research data.

- **Infrastructure Layer**: The foundation is built on Dell PowerEdge servers, which include hardware-level security features. The storage layer, utilizing Dell ECS or ObjectScale, provides essential data protection capabilities, including data encryption at rest, rich S3-compliant access policies, and object immutability options.
- **Network Layer**: The reference architecture for the DDL specifies a leaf-spine network design that provides traffic isolation and redundancy, forming a secure network fabric for all platform communications.
- **Platform Software Layer**: The DDL's system software, based on Kubernetes, enables the implementation of role-based access control (RBAC) for managing platform resources and services. The core data engines provide sophisticated access controls. The Dell Data Analytics Engine (powered by Starburst) integrates with enterprise identity providers like LDAP/AD and supports Single Sign-On (SSO). Critically, it offers fine-grained access control, including the ability to define permissions at the row and column level, as well as dynamic data masking to protect sensitive information.

- **Auditing and Logging**: All components of the DDL stack, from the storage systems to the query engine, generate detailed logs of all access and administrative events. This provides the comprehensive audit trail necessary to meet the logging and monitoring requirements specified in the RFP.

Integration Strategy

An effective governance framework for a multi-component solution like the one proposed for the DCC cannot be managed in silos. A security policy applied to one component must be consistently enforced across all others. Therefore, the proposed strategy is services-led, focused on creating a single, unified governance framework that spans the entire platform.

Dell Professional Services will work directly with the DCC's governance stakeholders to define a comprehensive set of roles, policies, and data access tiers as outlined in Section 12.0 of the RFP (e.g., Administrators, Data Developers, Data Users, Auditors). Once this framework is defined, our experts will implement and enforce it consistently across the integrated technology stack:

- **In the Dell Data Lakehouse**: The defined roles will be mapped to the DDL's RBAC system to control access to the underlying data, compute resources, and query engines. Fine-grained policies will be implemented in Starburst to restrict access to sensitive datasets or specific columns within a dataset.
- **In CKAN**: The same roles will be configured within CKAN's permission system to control who can view metadata, publish new datasets, edit existing ones, or manage community submissions.

This holistic approach ensures that a user's permissions are consistent throughout their entire journey on the platform. A researcher who is granted access to a restricted health dataset will have that permission reflected in the CKAN portal, the DDL query engine, and the audit logs. This mature, integrated approach to governance is essential for maintaining security and compliance in a complex data commons environment.

Response to Section 12.0: Governance Requirements

Our solution implements a comprehensive, hybrid governance model that combines centralized oversight with distributed stewardship, supported by a robust technical framework. This model is designed to be fully compliant with all legal and ethical requirements.

- **Core Capabilities (12.1)**:
  - **Dataset Auditing and Bias Detection**: The integrated AIF360 toolkit provides automated and manual auditing of datasets to detect biases across protected attributes. Statistical fairness metrics and visualizations will be used to flag datasets that pose a risk of reinforcing inequities.
  - **Model Fairness Evaluation**: The integrated Fairlearn and AIF360 libraries will be part of the standard ML development workflow, enabling users to assess and compare model performance across subpopulations and document the results.
  - **Bias Mitigation Tools**: The solution offers a comprehensive set of pre-processing, in-processing, and post-processing bias mitigation tools from the AIF360 library, allowing users to address identified fairness issues.

    - **Ethical Oversight and Review**: The platform will support a formal workflow for the ethical review of datasets and models, including annotations and approvals by designated reviewers, such as a Data Access & Use Committee.

- **Operational Requirements (12.2)**:
    - **Transparency and Auditability**: All fairness assessments, mitigation efforts, and governance decisions will be logged in the DDL's central audit trail. Public-facing summaries of fairness evaluations can be generated to foster transparency.
    - **Usability**: The fairness and bias tools will be made accessible through user-friendly Jupyter Notebook templates, providing recommended workflows and documentation for users who are not fairness experts.
- **Governance and Legal Compliance (12.3)**: Dell Professional Services will work with the DCC to ensure the platform's configuration and policies align with all state and federal anti-discrimination laws and emerging AI regulations. The governance structure will define clear roles for fairness auditors, data stewards, and reviewers.
- **Technical Integration (12.4)**: The fairness tools (AIF360, Fairlearn) will be fully integrated with the DDL's Spark-based Jupyter environment. APIs will allow for programmatic access to fairness functions, and the platform will support ongoing monitoring of deployed models for fairness drift.
- **Monitoring and Reporting (12.5)**: The system will be configured to automatically generate fairness reports for datasets and models. Dashboards will be provided to monitor fairness metrics across projects, with alerts triggered if models show signs of performance or fairness drift.
- **Education and Capacity Building (12.6)**: Dell Education Services will develop and provide training modules to help users understand fairness concepts, the integrated tools, and ethical considerations, building a community of responsible AI practitioners.

Response to Section 13.0: Security and Access Control Requirements

The proposed solution is built on a "security-first" architecture that provides defense-in-depth, from the physical infrastructure to the application layer.

- **Requirement 13.1 Access Controls**:
    - **User Authentication**: The platform will integrate with a central identity provider for Single Sign-On (SSO) and will enforce Multi-Factor Authentication (MFA) for access to sensitive data and administrative functions.
    - **Role-Based Access Control (RBAC)**: A comprehensive RBAC model will be implemented, defining roles such as Public Viewer, Researcher, Curator, and Admin. Access to restricted datasets will be granted only through a formal approval workflow managed within the platform.
    - **Dataset Permissions**: Permissions will be configurable at the dataset level, supporting tiers for Public, Restricted (requires review), and Internal access.
    - **Logging & Auditing**: The DDL provides a centralized and immutable audit trail that logs all data access, downloads, queries, and administrative actions. Logs will be retained for a minimum of 3 years to meet compliance requirements.

- **Requirement 13.2 Data Protection**:
    - **Encryption**: All data will be encrypted at rest on the Dell ECS storage platform and in transit using industry-standard protocols (currently: TLS 1.3).
    - **Threat Detection**: The platform will integrate Mass Tech Digital Loss Prevention (DLP) and Threat Vulnerability Management (TVM) intrusion detection and monitoring tools to protect against unauthorized access and cyber threats.
    - **Security Reviews**: Dell Managed Services will conduct regular security reviews and penetration testing to ensure the ongoing integrity and security of the platform.

Response to Section 14.0: Governance Policies and Ethical Guidelines

The platform is designed not only to be technically secure but also to enable and enforce the ethical guidelines and governance policies of the DCC.

- **Requirement 14.1 Guiding Principles**:
    - **Transparency**: The CKAN portal will provide public documentation on data sources, transformations, and usage licenses. The use of open-source components for fairness and credits further enhances transparency.
    - **Accountability**: The comprehensive audit trail tracks all data access and usage, enforcing responsible and accountable use of the DCC's resources.
    - **Equity**: The integrated AI Fairness and Bias Capability provides the tools to actively identify and mitigate the inclusion or use of biased data, promoting equity in AI development.
    - **Consent and Compliance**: The platform's governance framework will ensure that all datasets are managed in compliance with applicable laws and consent agreements.

Response to Section 2.0: Program Management and Service Delivery Excellence

Direct Mapping of Services to Dell Offerings

The RFP for the DCC clearly articulates the need for a full-service partner to provide not just technology but also comprehensive project management and technical consulting services. The Dell Professional Services portfolio is structured to provide exactly this end-to-end engagement, ensuring a successful and low-risk delivery of the DCC platform. Each service required in the RFP maps directly to an established, mature offering within the Dell Services portfolio, as detailed in the Dell Data Lakehouse technical documentation.

The following table illustrates this direct alignment:

RFP Required Service (Section 2.0)	Corresponding Dell Professional Services Offering	Description of Service
Program Management & Coordination	Dell Project Management Office (PMO)	Provides a dedicated Project Manager to develop and maintain the integrated project plan, conduct regular stakeholder meetings, manage risks, and ensure all milestones are met on time and within budget.
Technical Architecture & Platform Development	ProConsult Advisory for GenAI & Accelerator Services for Dell Data Lakehouse	Engages expert architects to design the detailed platform architecture, integrating the DDL with CKAN, custom-built Common Credit Solution, and other components. The Accelerator Services provide rapid deployment and configuration of the core DDL platform.
Dataset Curation and Integration	ProConsult Advisory Services for Data Preparation work	Provides data engineering expertise to assist the DCC in curating initial datasets, developing data ingestion pipelines, ensuring metadata completeness, and implementing versioning and provenance tagging and configuring connections of Starburst Galaxy Gravity (SGG) data catalog to the 24 data sources specified in the RFP and connecting the CKAN data catalog platform to the SGG data catalog that is part of the core Dell Data Lakehouse platform
Synthetic Data & AI Fairness Capability	ProConsult Advisory & Implementation Services	Dell's AI experts will architect and implement the integration of synthetic data and fairness tools (SDV, AIF360) into the DDL's Spark environment, including the development of template notebooks
Security, Governance & Compliance	Advisory Services for GenAI Data Security	Provides security and governance experts to work with the DCC to define and implement the unified RBAC model, data tiering, encryption, logging, and ethical use policies across the entire platform. Integrate DAF & DLH with Mass Tech Digital Loss Prevention (DLP) & Threat Vulnerability Management (TVM) platforms
Testing, Rollout, and Adoption	Dell QA Engineering & Dell Education Services	Conducts comprehensive functional, performance, and security testing. Dell Education Services will develop and deliver training for end-users and administrators and prepare all necessary

RFP Required Service (Section 2.0)	Corresponding Dell Professional Services Offering	Description of Service
		documentation and support materials for a successful public launch
Ongoing Support	ProSupport Infrastructure Suite & Advisory Services Subscription	Provides 24/7 technical support for the entire infrastructure stack. The advisory subscription offers a dedicated expert to help the DCC maximize value from the platform post-launch

The Dell Advantage

By selecting this proposal, the Commonwealth is procuring more than a technology stack; it is securing a strategic partnership with a single, globally recognized leader in enterprise data and AI platforms. The primary advantage of this approach is the mitigation of risk and the simplification of project governance. Instead of managing multiple vendors for hardware, software, and integration services, the DCC will have a single point of contact and accountability in Dell for the entire end-to-end solution.

This unified approach ensures that all components are designed to work together from day one, avoiding the common pitfalls of multi-vendor integration projects. Our proven methodologies, encapsulated in the Dell Professional Services portfolio, have been refined over thousands of successful enterprise data platform deployments. This experience ensures that the project will adhere to a structured, predictable, and transparent process, leading to a higher quality outcome delivered on schedule. The client is not just buying technology; they are buying a guaranteed outcome, backed by the full resources and expertise of Dell Technologies.

Providing the Services

Our proposed approach to delivering the Data Commons Collaborative is a comprehensive, services-led engagement managed by a single, accountable partner: Dell Technologies. We will leverage our proven project management methodologies and our portfolio of expert Professional Services to ensure the successful design, development, and launch of the DCC, meeting all required deliverables and milestones as defined in the project plan.

Our engagement is structured around the entire project lifecycle, from initial strategy to long-term operational success:

1. **Initiation & Design (Phase 1)**: The project will commence with a series of collaborative workshops led by our **ProConsult Advisory** architects. In this phase, we will finalize the detailed technical architecture, define the comprehensive governance and security framework, and approve the initial dataset curation criteria with all stakeholders. This phase concludes with the delivery of a Detailed Project Plan and Timeline and the Detailed Platform Architecture Documentation.
2. **Infrastructure & Platform Development (Phase 2)**: Our **ProDeploy for Infrastructure** and **Accelerator Services for Dell Data Lakehouse** teams will deploy and configure the core DDL

platform, including all compute, storage, and networking. Concurrently, our integration specialists will deploy and configure complementary open-source components, including the CKAN data portal, the Open Loyalty engine for the Commons Credits System, and the required AI fairness and synthetic data libraries. This phase delivers a fully operational, integrated platform.

3. **Capability Implementation & Testing (Phase 3)**: With the platform in place, our data and AI experts will work with the DCC team to ingest the initial curated datasets, implement the fairness and bias reporting dashboards, and generate sample outputs from the synthetic data tools. Our dedicated QA team will conduct rigorous functional, performance, and security testing across all features to ensure the platform is robust and ready for launch.

4. **Rollout & Adoption (Phase 4)**: **Dell Professional Services** and **Dell Education Services** will jointly lead the final phase, developing and delivering comprehensive training for both end-users (researchers, developers) and DCC administrators. We will finalize all user documentation and support materials. Dell will provide full support for the public launch and transition the platform to our **ProSupport** team for ongoing operational support, ensuring the long-term success and sustainability of the Data Commons Collaborative.

This structured, expert-led approach ensures a high-quality, on-time delivery that fully realizes the vision of the Massachusetts AI Hub.

Conclusion and Strategic Advantages

The proposed solution offers the Massachusetts AI Hub a unique and powerful combination of strengths, specifically architected to meet the diverse requirements of the Data Commons Collaborative. The strategic advantages of this approach are clear:

1. **Enterprise-Grade Foundation**: The solution is built upon the Dell Data Lakehouse, a secure, scalable, and high-performance platform engineered for the most demanding enterprise data and AI workloads. This provides the DCC with a reliable foundation that can grow with the needs of the Commonwealth for years to come, backed by the support and expertise of Dell Technologies.

2. **Open-Source Flexibility**: The platform strategically embraces the best of the open-source world, integrating leading platforms like CKAN, AIF360, and Fairlearn. This aligns with the DCC's mission of transparency and collaboration, avoids vendor lock-in, and allows the platform to leverage the innovation of vibrant global communities.

3. **Unified Governance and Security**: The solution implements a single, coherent security and governance model that is enforced consistently across all components including the Mass Tech Threat Vulnerability Management (TVM) and Digital Loss Prevention (DLP) platforms. This holistic approach, designed and implemented by Dell's security experts, ensures that data is protected, access is controlled, and ethical guidelines are upheld throughout the platform.

4. **Single-Partner Accountability**: By choosing Dell, the Commonwealth is not just procuring hardware and software; it is engaging a single, accountable partner for the entire end-to-end project. Dell Professional Services will manage all aspects of design, integration, deployment, and support, dramatically reducing project complexity and risk for the Massachusetts AI Hub.

5. **Cost Management**: There can be variability in the timeline for delivery of the custom ‘user access management’ and ‘Common Credits’ functionality. The initial price quoted accounts for worst case scenario of a twelve-month delivery. Delivery of this custom solution will be managed in

Agile sprints. If the delivery of that solution is less than twelve months, the scope will be adjusted and unused backlog will be applied to future projects or feature enhancements.

Alignment with DCC Mission

Ultimately, the proposed solution is more than a technical architecture; it is a capability designed to accelerate the DCC's core mission. It is a platform engineered to empower AI innovation and research by providing centralized, ethical, and secure access to a vast range of data. It is designed to foster collaboration between the public, private, and academic sectors by making data resources discoverable and easy to use. By combining the reliability of an enterprise leader with the spirit of the open-source community, this solution provides the extensible and sustainable foundation the Commonwealth needs to position itself as a global leader in ethical and innovative AI.

Application

Respondents are cautioned to read this RFP carefully and to conform to its requirements. Failure to comply with the requirements of this RFP may serve as grounds for rejection of an Application.

a. All Applications must be submitted electronically.

Read and understood.

b. Required Submissions- All Applications must include the items listed below:
- Application Cover Sheet (Attachment A)

  Please refer to Attachment A as the Cover Sheet of this response.

- Application, which shall include:

  - A description of the firm responding to the RFP (including descriptions of proposed subcontractors, if any) and the firm's qualifications to perform the Services.

    Dell Technologies is a Fortune 500 global leader recognized for its innovation in digital transformation and enterprise IT solutions. Serving commercial, government, and public sector clients worldwide, Dell provides a comprehensive portfolio—including IT infrastructure, cloud services, data storage, servers, end-user devices, and cybersecurity—through business segments such as Infrastructure Solutions, Client Solutions, and OEM Solutions.

    Dell brings proven qualifications in Generative AI by offering integrated solutions across infrastructure, validated designs, and professional managed services. The Dell AI Factory platform leverages partnerships with NVIDIA, AMD, Intel, Hugging Face, and top data providers to deliver high-performance, scalable AI deployments with strong security, compliance (including FIPS), and data governance. Notable enterprise clients—including Samsung SDS, City of Amarillo, Kennedy Miller Mitchell, and Northwestern Medicine—have achieved measurable innovation and operational improvements with Dell's GenAI solutions. Dell's ongoing investment in certifications, services, and rapid pilot programs makes it a trusted partner for secure and transformative Generative AI initiatives.

  - The proposed approach to providing the Services. Additionally, Respondents are invited to propose alternative(s) which provide substantially better or more cost-effective performance than achievable under the stated RFP scope of services.

    Please refer to the Services Required section of this response.

  - Resumes or bios of key personnel that will be performing the Services. – provide upon award

    Resumes or bios will be provided upon down selection.

  - Provide the total not-to-exceed costs for providing the Services based on projected hours, proposed hourly rates, as well as any other appropriate costs, by program phase in the Budget Template (Attachment C). List additional fees, overhead charges, or reimbursable expenses, if any. As a general policy, the Mass Tech Collaborative does not pay markups on reimbursables or out-of-pocket expenses. The Mass Tech

Collaborative also does not pay for word processing, overtime or meals. For travel costs, the Mass Tech Collaborative pays the IRS rate per mile.

Please refer to Attachment C – Budget Template sent as an additional attachment to this response

- Three examples of work previously performed by the Respondent that are substantially similar to the Services. Respondent shall also provide references including a contact person, email and phone number. M- provide upon down selection

  Resumes or bios will be provided upon down selection.

- Respondent's W-9.

  Please refer to Dell's W-9 included in this response.

- Authorized Application Signature and Acceptance Form (Attachment B). **By executing the Authorized Respondent's Signature and Acceptance Form and submitting a response to this RFP, Respondents certify that they (1) are in compliance with the terms, conditions and specifications contained in this RFP, (2) acknowledge and understand the procedures for handling materials submitted to the Mass Tech Collaborative as set forth in subsection d. below, (3) agree to be bound by those procedures, and (4) agree that the Mass Tech Collaborative shall not be liable under any circumstances for the disclosure of any materials submitted to the Mass Tech Collaborative pursuant to this RFP or upon the Respondent's selection.**

  Please refer to the Authorized Application Signature and Acceptance Form (Attachment B) included in this response.

c. Applications must be submitted to proposals@masstech.org (please include the RFP number in the subject heading).

   Read and understood.

d. Any and all responses, Applications, data, materials, information and documentation submitted to Mass Tech Collaborative in response to this RFP shall become Mass Tech Collaborative's property and shall be subject to public disclosure. As a public entity, the Mass Tech Collaborative is subject to the Massachusetts Public Records Law (set forth at Massachusetts General Laws Chapter 66). There are very limited and narrow exceptions to disclosure under the Public Records Law. If a Respondent wishes to have the Mass Tech Collaborative treat certain information or documentation as confidential, the Respondent must submit a written request to the Mass Tech Collaborative's General Counsel's office no later than 5:00 p.m. ten (10) business days prior to the required date of Application submission set forth in Section 3.2 below. The request must precisely identify the information and/or documentation that is the subject of the request and provide a detailed explanation supporting the application of the statutory exemption(s) from the public records cited by the Respondent. The General Counsel will issue a written determination within five (5) business days of receipt of the written request. If the General Counsel approves the request, the Respondent shall clearly label the relevant information and/or documentation as **"CONFIDENTIAL"** in the Application and **shall only include the confidential material in the hard copy of the Application**. Any statements in an Application reserving any confidentiality or privacy rights that is inconsistent with these requirements and procedures will be disregarded.

   Read and understood.

Attachment B

Massachusetts Technology Collaborative
Authorized Respondent's Signature and Acceptance Form

The undersigned is a duly authorized representative of the Respondent listed below. The Respondent has read and understands the RFP requirements. The Respondent acknowledges that all of the terms and conditions of the RFP are mandatory, and that Respondent's response is compliant with such requirements.

The Respondent understands that, if selected by the Mass Tech Collaborative, the Respondent and the Mass Tech Collaborative will execute an Agreement specifying the mutual requirements of participation.

Respondent agrees that the entire bid response will remain valid for sixty (60) days from receipt by the Mass Tech Collaborative.

I certify that Respondent is in compliance with all corporate filing requirements and State tax laws.

I certify that the statements made in this response to the RFP, including all attachments and exhibits, are true and correct to the best of my knowledge.

I acknowledge and understand the procedures for handling confidential materials submitted to the Mass Tech Collaborative as set forth in Section 3.1(d); I agree to be bound by those procedures and agree that the Mass Tech Collaborative shall not be liable under any circumstances for the disclosure of any materials submitted to the Mass Tech Collaborative pursuant to this RFP or upon the Respondent's selection.

Respondent: Dell Marketing, L.P.
(Printed Name of Respondent)

By: * Laura Valdés
(Signature of Authorized Representative)

Name: Laura Valdes
Title: Proposal Manager
Date: October 6, 2025

* Dell is submitting this proposal subject to the Dell Technologies Commercial Terms of Sale found at www.dell.com/cts, including warranty, returns and service provisions. Any other terms in this Massachusetts Technology Collaborative Request for Proposals for AI Hub Data Commons Collaborative Consulting Services RFP No. 2026-JAII-01 or on a resulting Purchase Order shall not be applicable.

Attachment C – Budget Template

Please refer to Attachment C – Budget Template sent as an additional attachment to this response.

Pricing Notes

Pricing, if included, may change in case of a general increase of the cost of manufacturing or logistics, such as industry-wide shortages of components, resources or materials, government-imposed tariffs or fees, or other factors beyond Dell Technologies' reasonable control or in the event of structural or significant changes in the way Dell develops its list pricing. Products may be discontinued or revised (including components thereto) at any time without notice. Should the initial proposed product(s) be discontinued. Dell Technologies reserves the right to re-negotiate pricing.

Form **W-9**
(Rev. March 2024)
Department of the Treasury
Internal Revenue Service

Request for Taxpayer Identification Number and Certification

Go to *www.irs.gov/FormW9* for instructions and the latest information.

Give form to the requester. Do not send to the IRS.

Before you begin. For guidance related to the purpose of Form W-9, see *Purpose of Form*, below.

Print or type. See ***Specific Instructions*** on page 3.

1 Name of entity/individual. An entry is required. (For a sole proprietor or disregarded entity, enter the owner's name on line 1, and enter the business/disregarded entity's name on line 2.)

Dell Marketing LP

2 Business name/disregarded entity name, if different from above.

3a Check the appropriate box for federal tax classification of the entity/individual whose name is entered on line 1. Check only **one** of the following seven boxes.

☐ Individual/sole proprietor ☐ C corporation ☐ S corporation ☑ Partnership ☐ Trust/estate

☐ LLC. Enter the tax classification (C = C corporation, S = S corporation, P = Partnership) . . . . .

Note: Check the "LLC" box above and, in the entry space, enter the appropriate code (C, S, or P) for the tax classification of the LLC, unless it is a disregarded entity. A disregarded entity should instead check the appropriate box for the tax classification of its owner.

☐ Other (see instructions)

3b If on line 3a you checked "Partnership" or "Trust/estate," or checked "LLC" and entered "P" as its tax classification, and you are providing this form to a partnership, trust, or estate in which you have an ownership interest, check this box if you have any foreign partners, owners, or beneficiaries. See instructions . . . . . ☐

4 Exemptions (codes apply only to certain entities, not individuals; see instructions on page 3):

Exempt payee code (if any)

Exemption from Foreign Account Tax Compliance Act (FATCA) reporting code (if any)

(Applies to accounts maintained outside the United States.)

5 Address (number, street, and apt. or suite no.). See instructions.

C/O Dell USA LP P.O. Box 643561

6 City, state, and ZIP code

Pittsburgh, PA 15264-3561

Requester's name and address (optional)

7 List account number(s) here (optional)

Part I Taxpayer Identification Number (TIN)

Enter your TIN in the appropriate box. The TIN provided must match the name given on line 1 to avoid backup withholding. For individuals, this is generally your social security number (SSN). However, for a resident alien, sole proprietor, or disregarded entity, see the instructions for Part I, later. For other entities, it is your employer identification number (EIN). If you do not have a number, see *How to get a TIN*, later.

Note: If the account is in more than one name, see the instructions for line 1. See also *What Name and Number To Give the Requester* for guidelines on whose number to enter.

Social security number

			–		–				

or

Employer identification number

7	4	–	2	6	1	6	8	0	5

Part II Certification

Under penalties of perjury, I certify that:

1. The number shown on this form is my correct taxpayer identification number (or I am waiting for a number to be issued to me); and
2. I am not subject to backup withholding because (a) I am exempt from backup withholding, or (b) I have not been notified by the Internal Revenue Service (IRS) that I am subject to backup withholding as a result of a failure to report all interest or dividends, or (c) the IRS has notified me that I am no longer subject to backup withholding; and
3. I am a U.S. citizen or other U.S. person (defined below); and
4. The FATCA code(s) entered on this form (if any) indicating that I am exempt from FATCA reporting is correct.

Certification instructions. You must cross out item 2 above if you have been notified by the IRS that you are currently subject to backup withholding because you have failed to report all interest and dividends on your tax return. For real estate transactions, item 2 does not apply. For mortgage interest paid, acquisition or abandonment of secured property, cancellation of debt, contributions to an individual retirement arrangement (IRA), and, generally, payments other than interest and dividends, you are not required to sign the certification, but you must provide your correct TIN. See the instructions for Part II, later.

Sign Here **Signature of U.S. person** 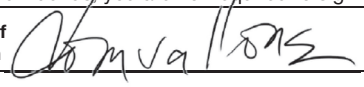 **Date** 1/6/2025

General Instructions

Section references are to the Internal Revenue Code unless otherwise noted.

Future developments. For the latest information about developments related to Form W-9 and its instructions, such as legislation enacted after they were published, go to *www.irs.gov/FormW9*.

What's New

Line 3a has been modified to clarify how a disregarded entity completes this line. An LLC that is a disregarded entity should check the appropriate box for the tax classification of its owner. Otherwise, it should check the "LLC" box and enter its appropriate tax classification.

New line 3b has been added to this form. A flow-through entity is required to complete this line to indicate that it has direct or indirect foreign partners, owners, or beneficiaries when it provides the Form W-9 to another flow-through entity in which it has an ownership interest. This change is intended to provide a flow-through entity with information regarding the status of its indirect foreign partners, owners, or beneficiaries, so that it can satisfy any applicable reporting requirements. For example, a partnership that has any indirect foreign partners may be required to complete Schedules K-2 and K-3. See the Partnership Instructions for Schedules K-2 and K-3 (Form 1065).

Purpose of Form

An individual or entity (Form W-9 requester) who is required to file an information return with the IRS is giving you this form because they

Proposal Legal Notes

Dell Technologies conducts operations through its subsidiaries and is the parent company to contracting legal entities Dell Marketing L.P. and EMC Corporation.

The contents of this response, including all elements of proposed pricing, performance level agreements and any referenced terms and conditions, apply only to direct purchases with Dell Technologies.

Terms & Conditions

This proposal will remain valid for 60 days from the date of submission of the proposal. Final pricing and other legally binding contract terms must be agreed or confirmed between the parties.

Dell is submitting this proposal subject to the Dell Technologies Commercial Terms of Sale found at www.dell.com/cts, including warranty, returns and service provisions. Any other terms in this Massachusetts Technology Collaborative Request for Proposals for AI Hub Data Commons Collaborative Consulting Services RFP No. 2026-JAII-01 or on a resulting Purchase Order shall not be applicable.

Disclaimer

This proposal (and information contained herein) is provided to you for information purposes only. Dell Technologies is not responsible for any errors or omissions relating to this proposal or that may occur as a result of the passage of time. In addition, Dell Technologies may improve or change this presentation or improve or change its products and service offerings from time to time, without updating this proposal. Please contact your sales representative for updates or additional information.

Confidentiality

This proposal (and information contained herein) is Dell Technologies Confidential Information, and your access and use are subject to and governed by the terms of your written nondisclosure agreement with Dell Technologies. In the absence of an applicable, written nondisclosure agreement between you and Dell Technologies, your access and use of this proposal (and information contained herein) shall be limited as follows: you will maintain the confidentiality of the Dell Technologies Confidential Information with at least the same degree of care that you use to protect your own confidential information, but no less than a reasonable degree of care under the circumstances; you may use the Dell Technologies Confidential Information only for the business transaction between you and Dell Technologies ("Purpose"); you may disclose Dell Technologies Confidential Information only to your employees who have a need to know the information for the Purpose and are legally bound by similar nondisclosure terms; and you will not disclose Dell Technologies Confidential Information to any other employee or to a third party.

Note

This information may be exempt from disclosure under open records and/or freedom of information act (foia) statutes and regulations. Dell reserves all rights available to it under applicable law to appeal any disclosure to a third-party accordingly.

Workstream		Services Cost	Expenses
1	**Consult**	**$2,695,571.20**	**$22,000.00**
1	SV-30643 - DATA PLATFORMS	$ 154,872.00	$ 13,200.00
2	SV-30652 - DEVSECOPS	$ 1,700,160.00	$ -
3	SV-30643 - DATA PLATFORMS	$ 202,504.00	$ 8,800.00
4	SV-30622 - TPO PROGRAM MANAGEMENT	$ 638,035.20	$ -
2	**AI Offerings**	**$ 217,562.40**	**$ 8,800.00**
1	SV-31091 - GENAI STRATEGY & GOVERNANCE	$ 175,982.40	$ 8,800.00
2	SV-31077 - GENAI MODELS & APPLICATIONS	$ 41,580.00	$ -
3	**Custom**	**$ -**	**$ -**
1	Describe Custom Service	$ -	$ -
Dell Services Total		**$ 2,913,133.60**	**$ 30,800.00**

Uncheck any items you want to hide:

- ☑ Custom Workstream
- ☑ Force Price
- ☑ Target Margin
- ☑ Total Force Price
- ☑ Total Force Margin
- ☐ Total Discount %

				Force Price	
Total Cost	List Price	List Margin		Sell Price	Sell Margin
$2,717,571.20	**$5,134,874.81**	**47.08%**		**$5,134,874.81**	**47.08%**
$ 168,072.00	$ 255,514.70	34.22%			
$ 1,700,160.00	$ 3,435,168.00	50.51%			
$ 211,304.00	$ 332,488.10	36.45%			
$ 638,035.20	$ 1,111,704.00	42.61%			
$ 226,362.40	**$ 395,784.67**	**42.81%**		**$ 395,784.67**	**42.81%**
$ 184,782.40	$ 281,736.67	34.41%			
$ 41,580.00	$ 114,048.00	63.54%			
$ -	**$ -**	**0.00%**		**$ -**	**0.00%**
$ -	$ -	0.00%			
$ 2,943,933.60	**$ 5,530,659.48**	**46.77%**		**$ 5,530,659.48**	**46.77%**

Target Margin		Total Force Price		Total Force
Sell Price	Sell Margin	Sell Price	Sell Margin	Sell Price
$ 4,180,878.77	35.00%	$ 4,183,776.45	35.05%	$4,180,278.02
$ 397,127.02	43.00%	$ 344,701.38	34.33%	$ 348,199.81
$ -	0.00%	$ -	0.00%	$ -
$ 4,578,005.79	35.69%	$ 4,528,477.83	34.99%	$ 4,528,477.83

Margin		
Sell Margin		Notes
34.99%		
34.99%		
34.99%		
34.99%		